

THE FUNDAMENTAL MULTI-BASELINE MODE-MIXING FOREGROUND IN 21 CM EOR OBSERVATIONS

BRYNA J. HAZELTON¹, MIGUEL F. MORALES¹, IAN S. SULLIVAN¹*Draft version March 1, 2013*

ABSTRACT

The primary challenge for experiments measuring the neutral hydrogen power spectrum from the Epoch of Reionization (EoR) are mode-mixing effects where foregrounds from very bright astrophysical sources interact with the instrument to contaminate the EoR signal. In this paper we identify a new type of mode-mixing that occurs when measurements from non-identical baselines are combined for increased power spectrum sensitivity. This multi-baseline effect dominates the mode-mixing power in our simulations and can contaminate the EoR window, an area in Fourier space previously identified to be relatively free of foreground power.

1. INTRODUCTION

Observing redshifted 21 cm neutral hydrogen emission from the Epoch of Reionization (EoR) is an exciting new tool for observational cosmology. These observations have the potential to reveal the timing and duration of reionization, to constrain models of star and galaxy formation and to determine what kinds of sources dominate reionization (for reviews see Morales & Wyithe 2010; Furlanetto et al. 2006). Several new arrays are under construction or have recently been built to observe the 21 cm power spectrum, including LOFAR (LOw Frequency ARray²), PAPER (Precision Array for Probing the Epoch of Reionization³) and the MWA (Murchison Widefield Array⁴).

The major challenge for EoR observations is the presence of astrophysical foregrounds that are 4-5 orders of magnitude brighter than the EoR signal. Early work recognized that the EoR signal could in principal be separated from the astrophysical foregrounds because the foregrounds have smooth spectra and so should be concentrated in the first few line-of-sight Fourier modes (k_{\parallel}) while the EoR signal would extend up to much higher k_{\parallel} modes (Morales & Hewitt 2004; Zaldarriaga et al. 2004; Jelić et al. 2008; Wang et al. 2006; Harker et al. 2009). Unfortunately this picture is somewhat complicated by the interaction of the foregrounds with realistic instruments, which have frequency dependent responses. This interaction is called mode-mixing because the chromatic instrumental response can add spectral structure to the astrophysical foregrounds, throwing foreground power to higher k_{\parallel} modes and obscuring the EoR signal. A number of recent papers have begun to shed light on these mode-mixing effects through detailed examinations of how foregrounds propagate through the instrument and analysis.

Precision mode-mixing simulations by Datta et al. (2010) first identified a distinctive wedge shape in \mathbf{k} space, with the mode-mixing predominantly below a $k_{\parallel} \propto k_{\perp}$ line. Investigations of the response of single baselines to flat-spectrum foregrounds by several groups

(Morales et al. 2012; Trott et al. 2012; Vedantham et al. 2012; Parsons et al. 2012) showed that the wedge shape is characteristic of smooth spectrum astrophysical sources interacting with the chromatic response of the baselines. This mode-mixing occurs because the baseline integrates over a region of the uv plane to form visibilities and the baseline length in wavelengths varies with frequency. This results in a slowly varying value for the visibility as the baseline moves through the uv space and this variation bleeds into the line-of-sight direction because of the physical size of the integration region (see Morales et al. 2012 for a more complete mathematical and pictorial explanation). The wedge shape comes about because longer baselines change length more quickly than shorter baselines, so the power at larger k_{\perp} is thrown into proportionally higher k_{\parallel} modes. These authors also identified a region called the “EoR window” at low k_{\perp} and high k_{\parallel} that should be relatively free of this kind of contamination.

While these single baseline effects are very important, the mode-mixed power in the Datta et al. (2010) simulations is dominated by a multi-baseline effect that we identify in this paper. Multi-baseline mode-mixing occurs when measurements from non-identical baselines are combined together to increase the power spectrum sensitivity, as they are for nearly all proposed observations. This new effect has a shape in \mathbf{k} space that is similar to the previously identified single baseline mode-mixing, with most of the mode-mixed foreground power thrown into the wedge, but it also contaminates the EoR window. The multi-baseline effect is present in all analyses that combine measurements from non-identical baselines or use images to make power spectra.

In the next section we introduce the multi-baseline mode-mixing mechanism using a simplified simulation with only one foreground source and show that it is fundamental to multi-baseline analyses. Then in section 3 we develop a precision simulation of the power spectrum with realistic foregrounds and we conclude in section 4.

2. MULTI-BASELINE MODE-MIXING

Visibilities from identical baselines (baselines with the same length and orientation) can be added coherently because they measure exactly the same angular modes on the sky. Coherent summing averages down the noise on the measurement, increasing the instrumental sensitivity

¹ University of Washington, Seattle, 98195² <http://www.lofar.org/>³ <http://astro.berkeley.edu/~dbacker/eor/>⁴ <http://www.mwatelescope.org/>

to the associated modes. Of course baselines that have very different lengths and orientations measure different angular sky modes and cannot be added coherently, but baselines that are similar but not identical are *partially coherent*, that is they measure some of the same modes but they are each sensitive to modes not detected by the other. In the uv plane, baselines integrate over a small region of the plane given by their power response, so the measured visibilities contain signals from a range of uv locations. For identical baselines the integration region is the same, but for partially coherent baselines the integration regions only partially overlap. Partially coherent visibilities contain separate measurements of the modes inside the overlap region, so it is common to coherently add the visibilities in the overlapping region to decrease the noise in those areas. All EoR sensitivity calculations explicitly combine partially coherent baselines (Morales & Hewitt 2004; Morales 2005; Bowman et al. 2006; McQuinn et al. 2006; Lidz et al. 2008; Beardsley et al. 2012b), as the alternative is a dramatic decrease in the potential sensitivity.

Unfortunately, combining measurements from non-identical baselines also introduces an additional type of mode-mixing. The reconstructed signal at a particular uv point contains co-added signal from an area of the uv plane given by the integration areas of all the co-added baselines. Because the locations of the baselines in the uv plane vary with frequency, the uv integration pattern for the reconstructed signal at the uv point also varies with frequency, introducing spectral ripple not present in the foregrounds themselves.

A visibility measurement is described by equation (12) of Morales & Matejek (2009)

$$\mathbf{m}([\mathbf{v}, f]) = \mathbf{B}([\mathbf{v}, f], [\mathbf{u}, f]) \mathbf{F}([\mathbf{u}, f], [\theta, f]) \mathbf{I}([\theta, f]) + \mathbf{n}([\mathbf{v}, f]) \quad (1)$$

which we have adapted to explicitly note the frequency dimension. This equation describes the integral that is being done by the instrument in the uv plane, the integral of the baseline power response (\mathbf{B}) multiplied by the Fourier transform of the true sky brightness ($\mathbf{F}\mathbf{I}$), and is mathematically identical to equation (1-13) in Clark (1999) up to the thermal noise term (\mathbf{n}). \mathbf{B} is only non-zero over a relatively compact region of the uv plane, so it defines the integration region for the baseline. It is important to note that neither the signal nor \mathbf{B} is flat over the integration region. The variation in \mathbf{B} over that region means that some areas of the uv plane contribute more strongly to the visibility than others and the integration over the signal variation leads to information loss on small uv scales (this is the source of the single-baseline mode-mixing described earlier).

For each uv location of interest, visibilities from all the overlapping baselines need to be added using a weighted sum. The reconstructed power at the uv location (\mathbf{u}_i) is given by

$$\tilde{\mathbf{I}}([\mathbf{u}_i, f]) = \tilde{\mathbf{B}}([\mathbf{u}_i, f], [\mathbf{v}, f])^T \mathbf{B}([\mathbf{v}, f], [\mathbf{u}, f]) \times \mathbf{F}([\mathbf{u}, f], [\theta, f]) \mathbf{I}([\theta, f]) \quad (2)$$

where $\tilde{\mathbf{B}}$ is the baseline weighting function (and we have dropped the noise term for clarity). $\tilde{\mathbf{B}}^T \mathbf{B}$ then records

how much signal from each uv location is mapped to the reconstructed power at \mathbf{u}_i . The movement of baselines in the uv plane with frequency generates significant frequency dependence in $\tilde{\mathbf{B}}^T \mathbf{B}$ because which baselines contribute to \mathbf{u}_i , and the locations of those baselines relative to \mathbf{u}_i , is strongly frequency dependent. This frequency dependence in the signal mapping generates frequency dependence in $\tilde{\mathbf{I}}$ even if the true sky brightness is spectrally smooth. Since the power spectra are generated by Fourier transforming in frequency along each \mathbf{u}_i , the spectral ripple introduced by combining partially coherent baselines throws power into high k_{\parallel} modes.

2.1. Multi-Baseline Simulation

To illustrate this effect, we developed precision simulations based on software holography/A projection (Morales & Matejek 2009; Bhatnagar et al. 2008) where the visibilities are gridded using the antenna response function ($\tilde{\mathbf{B}}$ is just a sampled version of \mathbf{B}). We use this approach because it can be shown to be information lossless (Tegmark 1997a,b), but the results are general to any analysis that combines measurements from non-identical baselines. Following Datta et al. (2010), we used the array layout from the originally proposed MWA instrument (Beardsley et al. 2012a) with 512 square antennas (Tingay et al. 2013).

Figure 1 demonstrates the effect of multi-baseline mode-mixing using a simplified foreground with a single flat-spectrum point source and focusing on one location in the uv plane. In the bottom panel the true phase of the source at the selected uv location is shown in black (referenced to the lowest frequency). The red line shows the reconstructed phase at that location as a function of frequency, exhibiting obvious frequency structure absent in the true foreground signal. We have selected three representative frequencies, indicated by the dashed blue lines, picking frequencies with almost no phase error, large negative and large positive error (left to right). Our selected uv point is in the lower left-hand quadrant of the uv plane and the upper nine panels zoom into that point, marked by a red cross, at each of the chosen frequencies.

In the top row, the integration regions for all the baselines that contribute to the uv point at any frequency are plotted in light grey. The baselines that contribute at the selected frequency are outlined in black, defining the region of the uv plane that contributes to the reconstructed signal at our point. The set of baselines that contribute change with frequency because the baseline length increases with frequency (baselines move to the lower left in subsequent columns) and the integration regions for the other columns are shown in middle grey.

The blue heat map shows normalized $\tilde{\mathbf{B}}^T \mathbf{B}$ or the relative influence of the areas within the integration region on the reconstructed signal at our uv point. Each visibility is formed by integrating over an area in the uv plane with the antenna response function and then the visibilities are weighted by the response at the uv point and averaged to determine the best estimate of the signal there. In other words, the reconstructed signal is a weighted average of the contributing visibilities. The influence map is the effective weight of each location within the integration region on the reconstructed signal.

The second row shows the true phase of our foreground

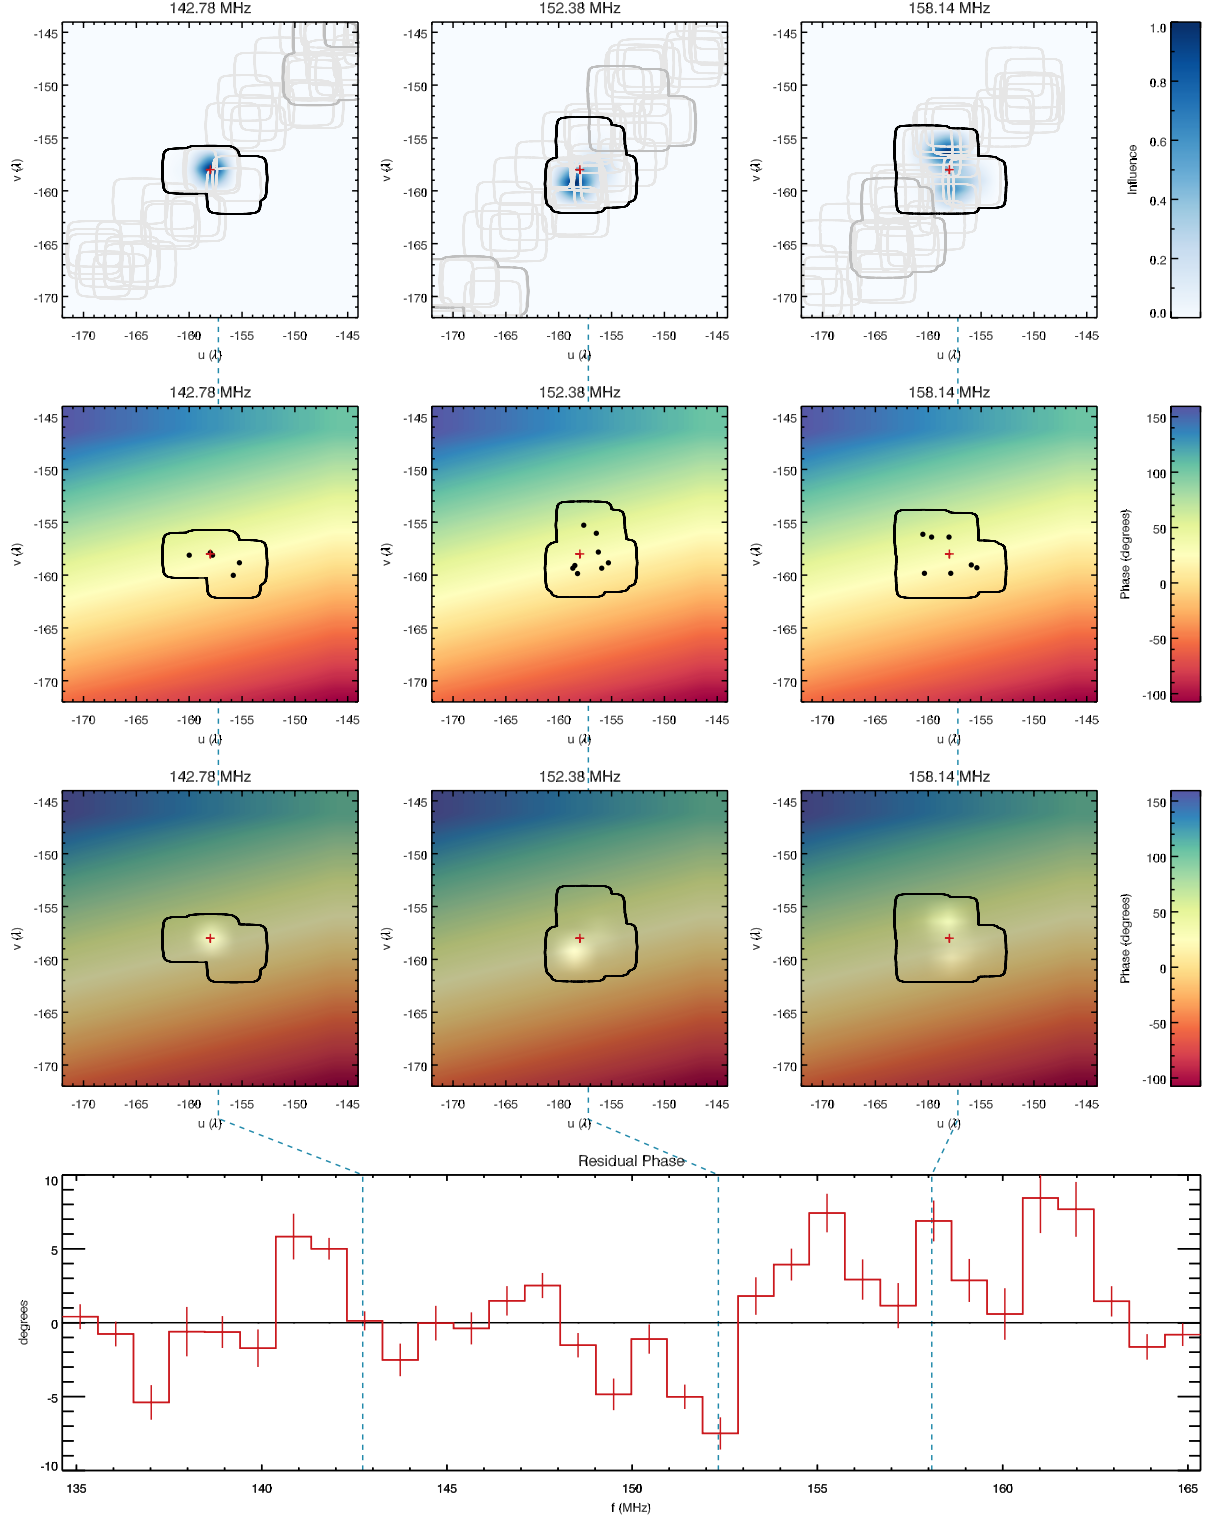


FIG. 1.— This figure demonstrates the effect of multi-baseline mode-mixing in the simplified case of one flat spectrum source at a single location in the uv plane. The *bottom panel* shows the true (flat black line) and reconstructed (red line) phase of the source (relative to the lowest frequency) at the uv location with the selected frequencies marked with blue dashed lines. The upper nine panels zoom into an area around the uv point at the three selected frequencies. The uv point is in the lower left-hand quadrant of the uv plane ($uv = 0$ is far to the upper right) and is marked by a red cross in each panel. The *top row* shows the integration areas for all the baselines that contribute to the uv point at any frequency in light grey and the set of baselines that contribute at the selected frequencies are outlined in black, defining the region that contributes to the uv point. The integration areas are given by the response functions of the square MWA antennas. Baseline lengths increase with frequency, moving to the lower left in subsequent columns, and the integration regions for the other columns are shown in middle grey. The blue heat map shows the relative influence of areas within the integration region on the reconstructed signal. The *second row* shows the true phase of the foreground source with the integration region and the centers of the contributing baselines marked in black. In the *third row* the influence map from the top row is overlaid on the true phase from the second row. This figure is discussed in detail in section 2.

source with the integration region and the centers of the contributing baselines marked in black (each dot corresponds to a grey square in the top row). Finally in the third row we have overlaid the influence map from the top row on the phase from the second row. For the left-hand column the phase is almost perfect because there are two visibilities that are nearly on top of our uv point, dominating the influence map, and the contributions from the other three nearly cancel out. In the center column most of the visibilities are at lower phase and the closest baselines, which dominate the influence map, are clustered just to the left and down from the uv point. In the third column the number of close visibilities above and below the uv point are nearly equal and the influence map is double-lobed but skewed slightly toward higher phase.

The variation in the reconstructed signal with frequency is caused by the dithering of which baselines contribute and how strongly they contribute (i.e. the influence map) as a function of frequency. The shape of the multi-baseline mode-mixing ripple is different in each uv pixel and depends on the details of the foregrounds and on the exact locations of the baselines contributing to each location, but unlike the mode-mixing discussed in previous work (Morales et al. 2012; Trott et al. 2012; Vedantham et al. 2012; Parsons et al. 2012), the spatial frequency (k_{\parallel}) of this ripple is not limited by the field of view of the instrument so it can throw power into higher k_{\parallel} modes, including into the EoR window.

2.2. Maybe we should....

While we have used gridding to show the origin of this effect graphically, multi-baseline mode-mixing is inherent in any approach that combines non-identical baselines to create power spectra. However the reasons for this are subtle and it is natural to suggest different analysis approaches to get around the problem. Common ideas include:

Maybe we should use a different gridding kernel. The effect of using a different gridding kernel in this case is simply to change the weighting of the different visibilities in the reconstructed signal. However the influence map is a combination of the integral done by the instrument to form the visibility and the gridding kernel, so it won't change dramatically and there will still be shifts in the map that will lead to mode-mixing. In addition, using the antenna response function as the gridding kernel is optimal (Tegmark 1997b) and using any other gridding kernel would result in a loss of sensitivity to the EoR signal.

Could we use different antenna elements to decrease this effect? The frequency dependence of $\tilde{\mathbf{B}}^T \mathbf{B}$ is dominated by the change in the baseline length as a function of frequency, not by the frequency dependence of the antenna response, so different antenna elements will not help much. The best way to mitigate this effect when designing an instrument is to have a smooth, dense baseline distribution, especially for small uv modes where the EoR signal is the strongest, because the accuracy of the reconstructed power is proportional to the density of the uv coverage.

What if we changed the visibility weighting to make the influence map frequency independent? It is not possible to weight the visibilities to achieve a frequency indepen-

dent influence map because most of the integration is done by the instrument to form the visibilities. For instance in the three maps shown in figure 1, the black outlines show the areas of the uv plane that were integrated by instrument. With only 5-8 baselines that overlap the uv point at each frequency, there is clearly no way to weight them to create identical influence maps.

What about calculating the signal at each uv point without gridding? This calculation would involve a weighted sum of the visibilities and suffers from the same mode-mixing as in the gridding approach because the locations of the contributing baselines vary with frequency. Indeed the gridding approach is mathematically identical to a weighted-sum of visibilities at each uv point.

What if we only calculated the power spectrum at the baseline centers? To measure the power spectrum we need a measurement at each frequency for every uv point we are using. Since the baselines move in uv with frequency, there are no uv locations that have a baseline precisely centered on them at every frequency, leaving no visibilities to use in the analysis. To increase the data we might consider including baselines that are nearly centered on the uv location, effectively setting a radius around the uv point. This scenario puts us right back in the mode-mixing regime, however, because there will be some variation in the locations of the baselines we choose with frequency. The down selection in baselines used in this case would also result in a huge decrease in sensitivity.

Can we design an array (or chose baselines) to avoid this kind of mode-mixing by having no partially coherent baselines? Even if there were no baselines that were partially coherent at any frequency, the baselines that contributed to the same location at different frequencies would have different offsets from the uv point, so the influence map would still vary with frequency. In fact our simulations show that the amount of contamination from mode-mixing is decreased in regions with densely spaced baselines and is much worse in areas with very few partially coherent baselines.

Combining a sparse redundant array and the delay spectrum, as pioneered by PAPER (Parsons et al. 2012), does avoid the multi-baseline mode-mixing effect because visibilities from non-identical baselines are never combined and the Fourier transform is along the baseline track through the uvf cube, not along the frequency axis. However, this approach also leads to a significantly smaller EoR window and, for the same collecting area, correspondingly lower EoR sensitivity.

The key issue for multi-baseline analyses using the full EoR window is that the instrument does a weighted integral in the uv plane that cannot be reversed and the only way to estimate the signal at any particular uv point is to use the measurements from nearby baselines. Combining these measurements with a weighted average based on the antenna response function (either directly or by gridding) will provide the highest EoR sensitivity, but any approach will suffer from mode-mixing due to the dithering of the contributing baselines.

3. POWER SPECTRA

In figure 2 we show the extent of the mode-mixing contamination in \mathbf{k} space with more realistic foregrounds. The foreground sources for this simulation are con-

constructed using the differential source counts from the 6th Cambridge survey at 151 MHz (Hales et al. 1988) with a flux density range of 0.1 – 1 Jy, resulting in approximately five thousand sources randomly located in a 30 x 30 deg² field. The flux range was chosen to represent the dominant sources that would remain after traditional image-based deconvolution of the brightest sources.

This power spectrum was constructed by Fourier transforming in the frequency direction and then calculating a three dimensional power spectrum estimator using weights derived from $\tilde{\mathbf{B}}^T \mathbf{B}$. A full description of this estimator will be described in an upcoming paper (Hazelton et al. in preparation), but briefly it accounts for the difference in sensitivity of the two Fourier components (i.e. sine and cosine) for each mode and marginalizes over the covariance between different modes. The power spectrum is reduced to the two dimensions shown here by doing a variance weighted average along annuli in the k_x and k_y directions.

This power spectrum is similar to those shown in Datta et al. (2010), with a clear wedge shape and a less contaminated EoR window (above the dashed line). The lowest k_{\parallel} mode shown is the flat-spectrum ($k_{\parallel} = 0$) mode, which would contain all the foreground power if there were no mode-mixing. Our simulations naturally include the single baseline effects identified by Morales et al. (2012), Trott et al. (2012), Vedantham et al. (2012) and Parsons et al. (2012) because we form visibilities as described in equation (1), but we find that the multi-baseline effects dominate the mode-mixed power throughout the uv plane. The contamination that is present in the EoR window is a signature of multi-baseline mode-mixing which can produce contamination to arbitrarily high k_{\parallel} .

Fortunately there are two factors that mitigate the multi-baseline contamination at low k_{\perp} , leaving the mode-mixed power in the EoR window lower than in the wedge. The first effect is that the shorter baselines that contribute at low k_{\perp} move more slowly in uv space with frequency so there is less severe mode-mixing because the distribution of baselines contributing to a particular uv location varies more slowly. In addition, simulations show that areas with denser and smoother uv coverage are less badly affected by this type of mode-mixing, so a dense central core of baselines can significantly reduce the amount of contamination at low k_{\perp} . A dense central core is also desirable for increasing the sensitivity of arrays to the EoR signal, so it is a common feature among most of the arrays currently being built to detect the EoR.

4. CONCLUSION

We have identified a new type of mode-mixing that occurs when measurements from non-identical baselines are combined coherently to increase instrumental sensitivity to the EoR power spectrum. Multi-baseline mode-mixing is present in all analyses that combine measurements from non-identical baselines or use images to make power spectra, and although the majority of the contamination is in the wedge it also contaminates the EoR window.

The amplitude of the contamination is significantly lower for regions with smooth, dense uv coverage. It is the dominant mode-mixing term in our simulations




FIG. 2.— Power spectrum showing the mode-mixing contamination in a simulation with approximately five thousand point sources between 0.1 – 1 Jy (see section 3 for details). The EoR window is above and to the left of the black dashed line.

using the originally proposed MWA array with 512 antennas, a layout with remarkably smooth uv coverage in its core (Beardsley et al. 2012a). While the amplitude of the multi-baseline mode-mixing effect depends on the specifics of each instrument and can only be determined through precision simulations, the dominance for the 512 antenna MWA suggests that it will likely be the primary mode-mixing effect for the first generation of EoR instruments that use multi-baseline analyses.

Now that we understand the sources of the mode-mixing foreground identified by Datta et al. (2010), the natural question becomes how can we mitigate them. In the next installment of our informal series on the mode-mixing foreground we will explore how to mitigate the multi-baseline mode-mixing foreground. Conceptually, the frequency structure of mode-mixing contamination has a characteristic shape at each \mathbf{u}_i point given by the arrangement of the baselines, the antenna response, and the local uv foreground. Unfortunately we cannot just fit for the mode-mixing because the frequency shape depends non-linearly on the foregrounds, but we can solve the problem iteratively given a good model of the instrument. Because we will only subtract shapes typical of the instrumental response, we will preferentially remove the mode-mixing terms while leaving as much of the EoR signal as possible. In spirit this approach approximates the covariance inversion of Liu & Tegmark (2011) and Dillon et al. (2012), while being computationally straightforward.

It is not yet clear whether experiments based on single-baseline power spectrum analyses (e.g. PAPER) with lower levels of mode-mixing and a smaller EoR window or those opting for the higher sensitivity multi-baseline approach with increased mode-mixing (MWA, LOFAR) will be better able to detect the EoR power spectrum. The outcomes of the currently planned experiments will

inform the next generation of Hydrogen Epoch of Reionization Arrays (HERA).

ACKNOWLEDGEMENTS

This work has been supported by the National Science Foundation Astronomy Division through CA-

REER award 0847753 and NSF Postdoctoral Fellowship 1003314, and by the University of Washington. We wish to particularly thank Danny Jacobs, Adam Beardsley and Aaron Parsons for helpful discussions and feedback.

REFERENCES

- Beardsley, A. P., Hazelton, B. J., Morales, M. F., et al. 2012a, *Monthly Notices of the Royal Astronomical Society*, 425, 1781, 1781
- Beardsley, A. P., Hazelton, B. J., Morales, M. F., et al. 2012b, *Monthly Notices of the Royal Astronomical Society: Letters*, L15, L15
- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, *Astronomy and Astrophysics*, 487, 419, 419
- Bowman, J. D., Morales, M. F., & Hewitt, J. N. 2006, *The Astrophysical Journal*, 638, 20, 20
- Clark, B. G. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 180, *Synthesis Imaging in Radio Astronomy II*, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley
- Datta, A., Bowman, J. D., & Carilli, C. L. 2010, *The Astrophysical Journal*, 724, 526, 526
- Dillon, J. S., Liu, A., & Tegmark, M. 2012, arXiv, astro-ph.CO, arXiv:1211.2232
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Physics Reports*, 433, 181, 181
- Hales, S. E. G., Baldwin, J. E., & Warner, P. J. 1988, *Monthly Notices of the Royal Astronomical Society*, 234, 919, 919
- Harker, G., Zaroubi, S., Bernardi, G., et al. 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 1138, 1138
- Jelić, V., Zaroubi, S., Labropoulos, P., et al. 2008, *Monthly Notices of the Royal Astronomical Society*, 389, 1319, 1319
- Lidz, A., Zahn, O., McQuinn, M., Zaldarriaga, M., & Hernquist, L. 2008, *The Astrophysical Journal*, 680, 962, 962
- Liu, A., & Tegmark, M. 2011, *Physical Review D (Particles and Fields)*, 83, 103006, 103006
- McQuinn, M., Zahn, O., Zaldarriaga, M., Hernquist, L., & Furlanetto, S. R. 2006, *The Astrophysical Journal*, 653, 815, 815
- Morales, M. F. 2005, *The Astrophysical Journal*, 619, 678, 678
- Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *The Astrophysical Journal*, 752, 137, 137
- Morales, M. F., & Hewitt, J. N. 2004, *The Astrophysical Journal*, 615, 7, 7
- Morales, M. F., & Matejek, M. 2009, *Monthly Notices of the Royal Astronomical Society*, 400, 1814, 1814
- Morales, M. F., & Wyithe, J. S. B. 2010, *Annual Review of Astronomy and Astrophysics*, 48, 127, 127
- Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. 2012, *The Astrophysical Journal*, 756, 165, 165
- Tegmark, M. 1997a, *The Astrophysical Journal Letters*, 480, L87, L87
- Tegmark, M. 1997b, *Physical Review D (Particles and Fields)*, 55, 5895, 5895
- Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, *Publications of the Astronomical Society of Australia*, 30
- Trott, C. M., Wayth, R. B., & Tingay, S. J. 2012, *The Astrophysical Journal*, 757, 101, 101
- Vedantham, H., Udaya Shankar, N., & Subrahmanyam, R. 2012, *The Astrophysical Journal*, 745, 176, 176
- Wang, X., Tegmark, M., Santos, M. G., & Knox, L. 2006, *The Astrophysical Journal*, 650, 529, 529
- Zaldarriaga, M., Furlanetto, S. R., & Hernquist, L. 2004, *The Astrophysical Journal*, 608, 622, 622